BlueGene/L

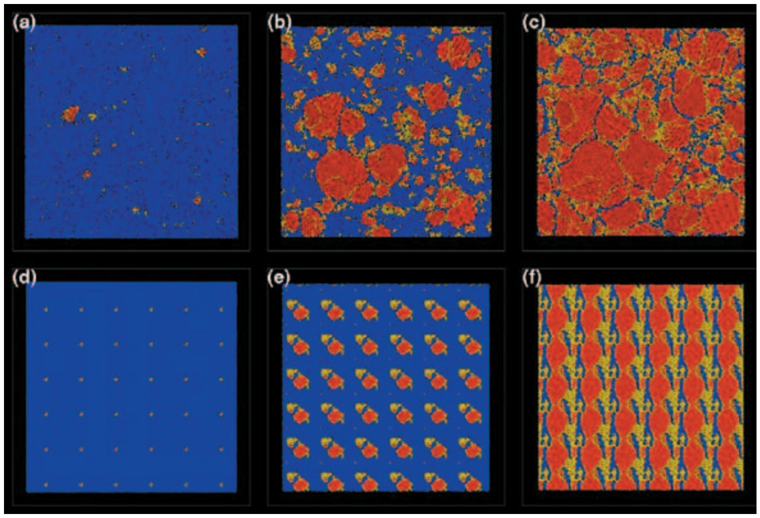

Snapshots from simulations of solidification in tantalum. The top sequence displays nucleation (a) and growth (b) occurring in a 16,372,000-atom simulation, resulting in a realistic distribution of grains and grain boundaries (c). The same process modeled using 64,000 atoms (d–f) produced the artificial final structure shown in (f).

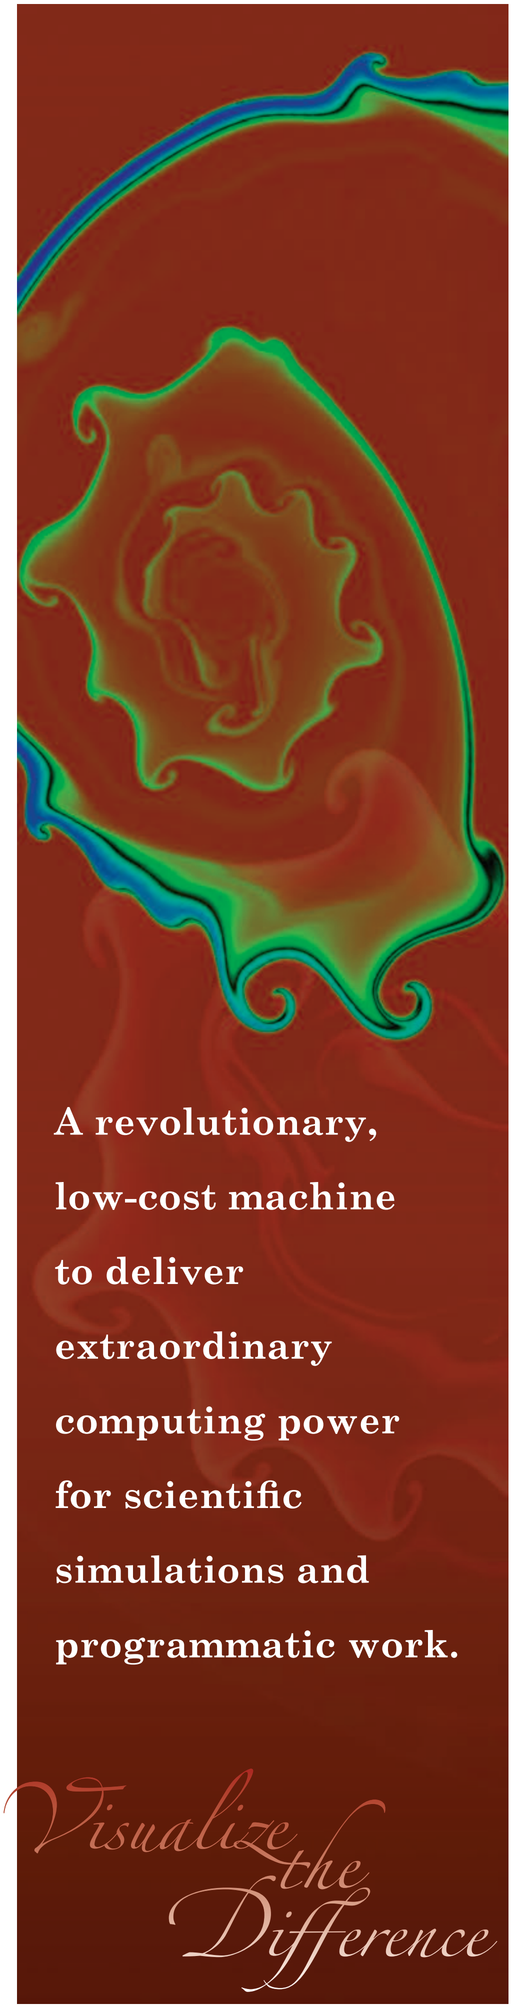

The IBM BlueGene/L supercomputer is fully assembled and is being prepared for classified service in support of the National Nuclear Security Administration's stockpile science mission of Los Alamos, Lawrence Livermore, and Sandia National Laboratories. The machine recently achieved a sustained world-record speed of 280.6 teraOPS on the Linpack benchmark at Lawrence Livermore, hinting at the machine's huge potential for rapid "time to solution" for applications in molecular dynamics and materials science.

Designed by IBM Research for the Department of Energy (DOE)/NNSA's Advanced Simulation and Computing (ASC) Program, the BlueGene/L supercomputer takes a radically different approach from supercomputers such as ASC Purple, Red Storm, and Q. Using a cell-based design, BlueGene/L is a scalable architecture in which the computational power of the machine can be expanded by adding more building blocks, without introduction of bottlenecks as the machine scales up.

The BlueGene/L computational research and evaluation platform development effort began with IBM in February 2000. The supercomputer uses system-on-a-chip (SOC) design technology and low-cost, low-power, embedded microprocessors. SOC design provides an extreme level of integration with all computing and network interfaces on a single custom application-specific integrated circuit only about 11 millimeters on a side. Low-power processors are a critical design element that enables BlueGene/L to achieve remarkable compute densities, providing 1,024 nodes (2,048 processors) and over 5 teraOPS in a single air-cooled cabinet.

BlueGene/L boasts a peak speed of over 360 teraOPS, a total memory of 32 tebibytes, total power of 1.5 megawatts, and machine floor space of 2,500 square feet. The full system has 65,536 dual-processor compute nodes. Multiple

ASC™

ADVANCED SIMULATION & COMPUTING™

Located in the Terascale Simulation Facility at Lawrence Livermore National Laboratory, BlueGene/L is optimized to run molecular dynamics applications at extreme speeds to address materials aging issues confronting the Stockpile Stewardship Program. BlueGene/L is also used to explore the potential of system-on-a-chip technologies to achieve extreme speed while minimizing floor space and electrical power consumption.

Programmatic research on BlueGene/L will provide the scientific insight to permit the development of increasingly predictive and accurate weapons codes.

communications networks enable extreme application scaling:

- Nodes are configured as a 32 x 32 x 64 3D torus; each node is connected in six different directions for nearest-neighbor communications
- A global reduction tree supports fast global operations such as global max/sum in a few microseconds over 65,536 nodes
- Multiple global barrier and interrupt networks allow fast synchronization of tasks across the entire machine within a few microseconds
- 1,024 gigabit-per-second links to a global parallel file system to support fast input/output to disk

The ASC tri-laboratory strategy for this cutting-edge computational platform is to significantly enhance material properties, turbulence, and shock-physics studies to improve the predictivity of simulations for the country's Stockpile Stewardship Program.

Scientific application time to solution will be improved by an order of magnitude on BlueGene/L. Key stockpile stewardship application results on the machine to date are pointing to a qualitative change in the way computational science can be performed. A number of target applications are running in excess of 50 teraOPS. With this rapid time to solution, scientists can perform a new run every day, make numerous investigations, and explore multiple alternatives. An entire scientific study can now be performed in the same time as just one science run required only a year ago. BlueGene/L is already making an impact on key NNSA missions and exploring one route toward cost-effective petaOPS computing capabilities.

UCRL-BR-216386

This work was performed under the auspices of the U.S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48

NNSA National Nuclear Security Administration

for more information

http://www.llnl.gov/NNSA/ASC